

PCT


WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04Q 11/04, H04L 12/56		A1	(11) International Publication Number: WO 95/24812
			(43) International Publication Date: 14 September 1995 (14.09.95)
(21) International Application Number: PCT/GB95/00502		(74) Agents: BLATCHFORD, William, Michael et al.; Withers & Rogers, 4 Dyer's Buildings, Holborn, London EC1N 2JT (GB).	
(22) International Filing Date: 8 March 1995 (08.03.95)			
(30) Priority Data: 94301673.3 9 March 1994 (09.03.94) EP		(81) Designated States: AM, AT, AU, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, JP, KE, KG, KP, KR, KZ, LK, LR, LT, LU, LV, MD, MG, MN, MW, MX, NL, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TT, UA, UG, US, UZ, VN, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG), ARIPO patent (KE, MW, SD, SZ, UG).	
(34) Countries for which the regional or international application was filed: AT et al.			
(60) Parent Application or Grant			
(63) Related by Continuation			
US		08/255,311 (CIP)	
Filed on		3 June 1994 (03.06.94)	
(71) Applicants (for all designated States except US): OXFORD BROOKES UNIVERSITY [GB/GB]; Gypsy Lane, Headington, Oxford OX3 0BP (GB). BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).		Published With international search report.	
(72) Inventors; and			
(75) Inventors/Applicants (for US only): SMITH, Avril [GB/GB]; 3 Finchdale Close, Carterton, Oxfordshire (GB). ADAMS, John, Leonard [GB/GB]; 24 Keswich Close, Felixstowe, Suffolk IP11 9NZ (GB).			

(54) Title: BROADBAND SWITCHING NETWORK



(57) Abstract

In a broadband switching system for the switching of asynchronously transferred cells of data, a dynamic bandwidth controller (DBC) controls the application of data cells to an input port of the system, the data cells being supplied by a number of transmitting end systems. The DBC provides buffering for cells from each source and controls the output of data cells to the system according to a current cell rate (CR) assigned to cells from the respective end system. When an end system begins transmitting data cells, the DBC detects the presence of incoming cells and requests bandwidth from a connection admission control (CAC) forming part of the system whilst providing a default cell rate for the end system so that transmission can continue until an appropriate bandwidth is allocated. The DBC contains a buffer for buffering cells in this situation. Once appropriate bandwidth has been allocated, cell rate advice information from the CAC is fed back to the DBC and thence to the transmitting end system to prevent the transmission of cells at too high a rate for the available bandwidth. The invention has both apparatus and method aspects.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

BROADBAND SWITCHING NETWORK

This invention relates to a broadband switching network for the switching of asynchronously transferred data cells, and
5 to a method of switching asynchronously transmitted data cells.

Broadband switching networks for switching asynchronously transferred cells are known, in which a predetermined level
10 of bandwidth is allocated to a transmission channel connecting a first customer to a second customer. In some of these known systems, a communications channel is provided over a significant period of time, effectively of the leased-line type, and manual measures are implemented in
15 order to set up such a connection or to modify a connection according to the particular terminations and the level of traffic being conveyed. Consequently, it is usual for customers to incur a fixed rate charge as part of the overall charge for the connection, resulting in payment
20 being made irrespective as to whether the connection is being used or not.

Alternative systems have been proposed or are available. In particular, it is possible for connections to be established
25 on a dial-up basis, requiring termination equipment to be provided with facilities for establishing connections by issuing signalling commands and responding to similar commands issued by the network.

30 The use of permanent circuits to support a private communications network is widespread. The demand for such circuits is expected to grow to include broadband rates above 2 Mbit/s, the circuits carrying traffic multiplexed from sources which are inherently bursty, possibly together
35 with traffic which is transmitted at constant bit rates and is delay sensitive, such as voice transmission and constant bit-rate video.

Asynchronous transfer mode (ATM) cells all have a fixed information field of forty eight octets which can carry customer traffic or customer-originating control information (signalling). These two types of data transmission are distinguished by setting virtual path (VP) and vertical circuit (VC) values in the cell headers. Another field provided in the ATM header is known as cell loss priority, which enables low priority cells to be distinguished from high priority cells. In the event of congestion, the low priority cells may be discarded first.

For private circuits within an ATM based network, the desired route, the required bandwidth, and the quality of service (QOS) are set up using network management procedures. The private circuits are known as permanent virtual circuits (PVCs) because there is no actual physical circuit, only a VP/VC value or "label" which is associated with information stored in the switches to determine the route and preserve the bandwidth and QOS requirements.

A disadvantage of all known permanent circuits is that the bandwidth remains assigned to the circuit, even when the customer has nothing to transmit. This means that the customer may have to pay higher charges than would be obtained if the bandwidth was only made available when needed. The assumption being made here is that charging is related to reserved bandwidth, and this is not necessarily correct in terms of the way public network operators may choose to charge for virtual circuits. However, it is expected that charging based on reserved bandwidth will become a significant factor in the future.

A common practice is to set up a permanent virtual circuit so that it is only available during certain hours of the day, or during certain days of the week. A difficulty with this approach is that it does not allow the customer to

change the pattern of usage quickly, and it may only crudely reflect the usage required by the customer.

A second proposal has been to provide the customer with a separate communications channel to the network management plane, thereby allowing a permanent virtual circuit to be reconfigured. A difficulty with this approach is that some time delay will be incurred before the customer can start to use the virtual circuit.

10

A third proposal is to introduce equipment at every switching point in the network that recognises a fast resource management cell, indicating that bandwidth should now be assigned to the circuit. A difficulty with this approach is that there is no internationally agreed standard for a bandwidth-requesting cell that would be recognised by the switching equipment produced by the various manufacturers.

20 According to a first aspect of this invention, a broadband switching system having at least one ingress for connection to a respective signal source and at least one egress for connection to a selected signal receiving system has at least one switch for transmitting information-carrying
25 asynchronously transferred data cells from the ingress to the egress, system control means for accepting and establishing a connection between the ingress and the egress via the said switch, and bandwidth control means arranged to detect the said information-carrying data cells received at
30 the ingress and, automatically in response to such cell detection, to cause the system control means to allocate bandwidth for the transmission of the cells to the egress. In this way it is possible to alter the bandwidth available in the system for cells intended for a given path
35 dynamically in response to the detected usage level, avoiding the need to signal a request for bandwidth prior to being able to transmit cells onto the system. It is also

possible to avoid the situation of bandwidth remaining assigned to a virtual circuit even when the customer has nothing to transmit.

5 Preferably, the bandwidth control means include a transmission limiter having storage means for storing a maximum output cell rate value derived from a control signal received by the limiter from the system control means, and means for limiting the transmission of cells to the said
10 switch when the rate at which cells are received at the ingress exceeds the maximum output cell rate value. The limiter may include a buffer for delaying transmission of the cells to the switch.

15 The bandwidth control means may include feedback means arranged to transmit the maximum output cell rate (CR) value to the ingress for transmission to a signal source generating the data cells. The maximum output cell rate value indicates a cell rate which depends on or matches the
20 allocated bandwidth and preferably forms part of a resource management (RM) cell. Thus, the rate at which data is fed to the ingress by the source can be controlled in a situation in which the bandwidth allocated depends on the bandwidth available in the switching system.

25

In the case of the bandwidth control means operating on data cells from a number of signal sources, the transmission limiter may be operable to transmit cells from one of the sources at a rate which is dependent on the number of cells
30 stored by the device and received from one or more of the other signal sources. Thus, the bandwidth available to signals outputted from the bandwidth control means may be shared between cells from different signal sources, especially if several transmissions having different route
35 identifiers are bursting simultaneously. This can be achieved by flagging cells which have been buffered for

longer than a predetermined period so as to give such cells a higher priority in a polling operation.

The bandwidth control means may include, in addition, a cell
5 monitoring stage operable to read routing identifiers associated with cells arriving from different signal sources and to update activity status values associated with the routing identifiers, the transmission of an activity status value relating to a particular identifier from an active to
10 an inactive status causing the system control means to allocate bandwidth for cells associated with that identifier.

Preferably, the bandwidth control means are arranged to
15 detect the rate at which cells are supplied to the ingress and bandwidth is requested from the system control means generally on the basis that the transmitting signal source should be allocated as much bandwidth as is available in the system for the indicated route. Such bandwidth will be
20 allocated depending on the level of traffic in the system at the time the system control means is caused to allocate bandwidth. However, in the preferred embodiment of the invention, the rate at which cells are supplied to the ingress from a given signal source is monitored and, if the
25 allocated bandwidth is more than that required to transmit the cells at the monitored rate, the system control means are caused by the bandwidth control means to allocate less bandwidth to these cells, e.g. after a predetermined time interval, thereby making bandwidth available to other
30 customers.

The bandwidth control means may also be arranged to read a bandwidth-identifying portion of an incoming cell, and to cause the system control means to allocate bandwidth to the
35 cell according to the identified bandwidth. The bandwidth control means may also be arranged to limit the rate of which incoming cells are accepted onto the network on the

basis of a default bandwidth until allocation of bandwidth by the system control means has occurred.

Buffering which may be provided in the bandwidth control means for buffering cells while bandwidth is being allocated may be used on other occasions when a signal source is transmitting at a higher rate than the rate capable of being accepted by the network at a given time. Indeed, it is preferable for the buffer to have means for detecting when it is filled to a predetermined threshold level, the feedback means being responsive to the buffer to cause a retransmission of a cell rate indicator signal to the signal source when connected to the ingress. Preferably, detection of filling of the buffer to a predetermined threshold level causes a reduced cell rate indicator signal to be transmitted to the signal source, the signal source having the facility then to reduce its cell rate to avoid buffer overflow and consequent loss of data.

The bandwidth control means may include a cell counter for counting cells received from respective signal sources coupled to the bandwidth control means, thereby to generate charging signals for customer billing and for other purposes. This counter may be part of an activity detector associated with a controller in the bandwidth control means, the activity detector being operable to detect incoming cells and identify the signal source transmitting the detected cells, with the controller being operable in response to the detection and identification to transmit a bandwidth request message to the system control means to allocate a level of bandwidth specifically to the cells transmitted by the identified signal source.

The system control means may include means for identifying available bandwidth in the system for the said allocation of bandwidth, the bandwidth control means including means operable to receive from the identifying means a control

signal representing the available bandwidth. It is possible for the system to re-allocate bandwidth from other existing connections in the system when the bandwidth control means requests bandwidth to be made available. Allocated
5 bandwidth may be reduced during transmission of cells by congestion being sensed and a congestion signal being transmitted to the bandwidth control means which, in turn, causes a feedback signal or "reduce traffic level" instruction to be sent to the ingress for causing a
10 reduction in the input signal rate from a respective signal source.

According to another aspect of the invention, there is provided a method of operating a broadband switching system
15 for transmitting information-carrying asynchronously transferred data cells from an ingress of the system to an egress of the system via at least one switch, the method comprising controlling the bandwidth for transmitting the cells by detecting the cells and, automatically in response
20 to the detection, causing system control means of the system to allocate bandwidth for transmission of the cells from the ingress to the egress.

25 According to a further aspect, the invention includes the use, in a broadband switching system for the transmission of information-carrying asynchronously transferred data cells from an ingress of the system to an egress of the system, of bandwidth control means arranged to detect incoming cells
30 received at the ingress and, automatically in response to such cell detection, to cause system control means of the system to allocate bandwidth for the transmission of the cells to the egress.

35 According to yet a further aspect of the invention, there is provided a dynamic bandwidth controller for a broadband switching system which is for transmitting information-

- carrying asynchronously transferred data cells from an ingress of the system to an egress of the system via at least one switch, wherein the controller comprises means for detecting the said cells received at the ingress, and means
5 for issuing a bandwidth request signal automatically in response to such cell detection to cause bandwidth to be allocated in the system for the transmission of the cells to the egress.
- 10 It is known to provide usage control means or so-called usage parameter controls (UPC) to downgrade the priority of cells or to discard cells according to priority designations as a means of controlling the application of traffic from an end system to the network. Where such usage control means
15 are present, the bandwidth control means may provide signals for the usage control means in response to the allocation of bandwidth by increasing a first predetermined threshold level at which the usage control means downgrades the priority of transmitted cells.
- 20 In the preferred embodiment of the invention, the first predetermined level of the usage control means is set to zero when no actual use is being made of the network by the end system, resulting in initial cells being downgraded to
25 low priority. Indeed, after a time-out period of the bandwidth control means having been exceeded with a particular end system having transmitted no data cells, the bandwidth allocated to that end system is reduced to zero.
- 30 Preferably, the bandwidth control means adjusts the first predetermined level set by the usage control means by transmitting a message to the system control means, which in turn issues a message to the usage control means.
- 35 The present invention also provides a method of asynchronously transmitting cells of data through a switchable network, the method comprising the steps of

resetting a usage control means so as to downgrade the priority of input cells; detecting the presence of input cells requiring a level of bandwidth through the network; and requesting a switch control to allocate bandwidth in response to said detection; whereupon bandwidth is allocated by the switch control; and the usage control means is adjusted by the switch control so that high priority cells are maintained as such.

10 In a preferred embodiment, a request may be made to the switch control to provide as much bandwidth as is available (typically up to the maximum transmission rate capacity of the transmitting end system), whereupon the switch control identifies the available bandwidth; the bandwidth is
15 allocated and buffer space is provided, and a feedback signal is generated, which is repeated if the level of data retained in said buffer exceeds a predetermined value.

The invention will now be described in greater detail by way
20 of example with reference to the accompanying drawings in which:

Figure 1 is a diagram of a broadband switching network in accordance with the invention;

25

Figure 2 is a diagram of another broadband switching network in accordance with the invention;

Figure 3 is a diagram of part of a broadband switching
30 network showing how a single bandwidth controller can be shared by several end systems;

Figure 4 is a block diagram of a bandwidth controller for use in the networks of Figures 1 and 2;

35

Figure 5 is a specification description language diagram (SDL) for the activity detector module shown in Figure 4;

10

Figures 6-1 and 6-2 are SDLs for the controller module of Figure 4;

Figure 7 is a diagram of a resource management (RM) data
5 cell;

Figure 8 is an SDL for the feedback module of Figure 4;

Figure 9 is a diagram of a buffer for the bandwidth
10 controller of Figure 4;

Figure 10 is an SDL for the buffer;

Figure 11 is a block diagram of a shaper/multiplexer module
15 and its connection to the buffer of Figure 6;

Figures 12-1, 12-2, and 12-3 are SDLs for the shaper/multiplexer module;

20 Figure 13 is a block diagram of a sub-switching network of an alternative broadband switching network in accordance with the invention, including a broadband asynchronous switch, a usage parameter control device and a dynamic bandwidth controller;

25 Figure 14 is a block diagram of the dynamic bandwidth controller of Figure 13, including a cell monitor; and

Figure 15 is a block diagram of the cell monitor of Figure
30 14.

In its preferred form, the invention is concerned with a broadband switching network which may form part of or may constitute a public switching network for the transmission
35 of asynchronously transferred data cells between end systems.

Referring to Figure 1, the public network 10 has a plurality of switches operable in asynchronous transfer mode (ATM). In this simple example, the switches include two local switches 12 each having a port for connection to a
5 respective end system 14, and a transit switch 16 interconnecting the local switches 12. Associated with the switches is a connection admission control function (CAC) 18 and a dynamic bandwidth controller (DBC) 20 for controlling traffic entering the network through one of the local
10 switches 12. This switch 12 also includes a usage parameter control device 22 for dynamically altering the priority of data cells received at the input port 24 of the network from the end system 14.

15 It will be understood that, in practice, the network 10 will include large numbers of local and transit switches 12, 16 and several DBCs 20 all interconnected to form a network having a plurality of ports such as port 24 for connecting several end systems such as end system 14. Using the DBC
20 20, the public network 10 is able to provide an available bit rate (ABR) service, the DBC acting to detect incoming cells supplied to the input port 24 and, automatically in response to this detection, to cause the CAC 18 to allocate bandwidth for the transmission of the cells to the
25 destination end system. Generally, end systems 14 requiring the ABR service are allocated to a fixed DBC 20. There may be more than one DBC 20 for each local ATM switch 12. In the case of a fault, end systems can be rerouted to a standby DBC (not shown).

30

Data is transmitted in the form of asynchronous transfer mode (ATM) cells, each having an information field of forty-eight octets, in addition to a header of five octets, which includes information facilitating transmission through the
35 network itself. Thus, routing is controlled on a cell-by-cell basis and a plurality of transmission paths and time multiplexed slots may be employed for any particular link.

ATM cells are, therefore, transmitted via virtual paths and virtual circuits, as defined by the header information.

The virtual paths and virtual circuits are identified by a
5 virtual path identifier (VPI) and a virtual channel
identifier (VCI) in the five octet header which effectively
defines the connection between the end systems so that cells
forming part of a common message will be transmitted over
the same connection. ABR traffic enters the public network
10 by routing cells according to their VPIs and VCIs through
the DBC 20 and then out to external routes, as shown in
Figure 1. From the DBC 20, the traffic on each virtual path
and virtual channel is restricted to a cell rate (which will
be referred to hereinafter as "CR"), determined by the CAC
15 18.

An alternative illustrative arrangement is shown in Figure
2. In this case, end system 14A is subject to the control
of more than one DBC. In fact, the connection between two
20 end systems 14A, 14B is routed through two public networks
10-1 and 10-2. Each network 10-1, 10-2 has its own DBC 20-
1, 20-2 responsible for restricting traffic entering the
network according to the bandwidth allocated by its own
connection admission control function (CAC) 18-1, 18-2.
25 Each DBC 20-1, 20-2 is also responsible for advising the end
system 14A of the current applicable CR.

In the systems of both Figure 1 and Figure 2, the DBCs 20,
20-1, 20-2 request bandwidth from the respective CAC 18, 18-
30 1, 18-2 whilst buffering any incoming data cells which
cannot immediately be transmitted to the respective switches
12, 16. The CAC 18, 18-1, 18-2 then allocates a bandwidth to
the cells corresponding to the most narrowband circuit in
the path identified by the VPI and VCI of the relevant
35 cells. This allocation is then indicated to the DBC 20, 20-
1, 20-2 which communicates a maximum CR to the transmitting
end system. As a consequence, the impression is given that

the network is permanently configured to make bandwidth available for data transmission, even when the user does not have a requirement to make use of the available bandwidth. However, as will become apparent hereinafter, the user need
5 only be charged when making actual use of the network, preferably by charging on the basis of the number of cells transmitted over the logical link between end systems. Thus, prior to transmission being invoked by, say, user A, the network may actually allocate no bandwidth whatsoever
10 for transmission between the relevant end systems. Furthermore, it is not necessary for any special control activity to be made by user A in order for bandwidth to be allocated. A level of bandwidth will be assigned to user A when required.

15

It is possible for a single dynamic bandwidth controller (DBC) to be shared by several end systems or signal sources. For example, referring to Figure 3, a DBC 20-3 is shown connected to a broadband ATM switch 12-3 forming part of the
20 network 10, the traffic of three sources 14C being handled using an output buffer 28. The number of sources which can be handled by the DBC 20-3 is determined by the link rate L (i.e. there must be not so many sources that it is always the link rate L which is the limiting factor determining the
25 available rate.) The aggregate cell rate of the ABR traffic from the sources 14C must not exceed L. This implies that if the traffic from each source is bursty, there may be times when the output buffer 28 is congested. This can be avoided by supplementing the sustained cell rate (CR)
30 feedback to the end systems 14C with generic flow control (GFC) signals which operate to stop all transmissions from each source immediately.

Whenever the dynamic bandwidth controller (DBC) is
35 incorporated in the arrangements of Figures 1, 2, or 3, its main functions are as follows.

Firstly, it provides buffering of incoming data cells, the degree of buffering at any given time being determined according to the transmission containing the cells, the transmission being identified by the VPI and VCI information referred to above. The DBC further controls or "shapes" the traffic fed to the network 10 so as to be equal to the current CR applicable to that particular transmission, the CR depending on the allocated bandwidth.

10 The allocated bandwidth, and hence the CR, for any given transmission is determined by the CAC 18 (see Figure 1) on the basis of determining the route to be followed by the transmission and assigning a fair share of the available capacity on the route based upon the known number of active
15 transmissions.

When a transmission begins it is detected in the DBC, which immediately transmits a low default CR to the relevant end system 14 (see Figure 1). This ensures that a newly active
20 transmitting source does not cause overload in the network 10 before the CAC 18 has been able to allocate bandwidth and derive an CR for that transmission. The traffic associated with that transmission leaving the DBC 20 is shaped to correspond to the default CR. This is part of a second main
25 function of the DBC, i.e. to send a feedback signal to the end system for the purpose of controlling its transmitted cell rate. Indeed, each time the CAC 18 derives a new CR for a transmission, a CR advice signal is fed back to the end system.

30 The pre-transmission buffering by the DBC is used to allow a cooperating end system sufficient time to adjust its output to the latest CR feedback advice. This implies that whilst the most recent CR is immediately used in the DBC to
35 shape output traffic from the DBC, there is sufficient buffering in the DBC to allow excess cells to enter for a period at least equal to the round trip delay between the

15

DBC and the end system. If cells continue to arrive from the end system 14 at a rate greater than the advised feedback CR (for instance, because the CR was lost en route, or because of a faulty end system) the excess cells
5 will be dropped in the DBC by overflow of the buffer.

In the preferred DBC, it is also possible to include fault tolerance by making use of a buffer threshold. When the stored cells relating to a given transmission reach the
10 threshold, retransmission of the CR advice feedback to the end system is triggered. This feature is useful also as a mechanism for policing end systems to prevent inefficient use of bandwidth, whether due to a faulty terminal or due to deliberate non-compliance with contracted transmission
15 rules. In this way, interference with the quality of service provided for other, compliant end systems is prevented. In effect, the DBC defines the ABR traffic contract with the network 10.

20 The modules of the DBC 20 will now be described in more detail with reference to Figure 4.

The DBC 20 is shown in Figure 4 is a discrete unit having an input port 30 for receiving asynchronously transmitted data
25 cells, an output for feeding data cells onto a switch 12 or 16 (see Figures 1 and 2) forming part of the switching network 10. The unit also has another input 34 for receiving messages back from the switch 12 or 16 and a feedback output 35 for transmitting feedback messages to the
30 end system 14 (shown in Figure 1). Although the DBC 20 is shown as a discrete unit, it will be appreciated that Figure 4 can be regarded as a functional diagram representing a subset system of a larger data processing unit, much of which may be embodied as software functions.

35

Incoming cells on input 30 arrive as a user cell stream which is fed firstly to an activity detector 36. The

purpose of the activity detector is to provide state information to a controller module 38 about each received transmission, each transmission being identified by its VPI and VCI contained in the cell headers. A transmission is
5 labelled active by the activity detector 36 if it was previously quiet and a cell having the appropriate VPI and VCI values is observed to be transmitted from an end system to input 30. Synchronisation of the activity detector 36 with the start of a cell header may be carried out using an
10 error check field contained in the cell header. A transmission is considered to be in an inactive state if it was previously active and no cell having the appropriate VPI and VCI values has been detected for a period of time t . Activity detector 36 maintains a timer and state table for
15 each VPI/VCI value pair, ready for an interrogation by the controller 38. Preferably, t is set to be several seconds so that active-inactive-active transitions relating to any given VPI/VCI value pair which are of the order of several milliseconds remain undetected so that the transmission is
20 indicated as remaining in the active state under these conditions. This has the effect of reducing the frequency of messages sent by the DBC 20 to the connection admission control CAC 18 at some expense to lowering utilisation of the network.

25

Another function of the activity detector 36 is that of counting the cells for a transmission during an interval between a "get cell count" and a "stop cell count" request from the controller 38. This information can be used, for
30 example, for charging purposes and also by the controller 38 for assessing the actual cell rate of received transmissions.

Pseudocode for the activity detector is listed below and the
35 corresponding SDL appears in Figure 5.

```
BEGIN {cell arrival}
    cell arrival from end system
```

17

```

        read VC
        reset VC inactivity timer
        IF VC is newly active THEN
            update state table
5         advise CONTROLLER of newly active VC
        ELSE IF counting.cells(VPI/VCI) THEN
            increment cell.count(VPI/VCI)
        ELSE
            do nothing
10      END

        BEGIN {VC timer expires}
            VC inactivity timer expires indicating quiet VC
            update state table
15         advise CONTROLLER of quiet VC
            counting.cells :=FALSE
        END

        BEGIN {receive a start.cell.count signal}
20         receive a start.cell.count(VPI/VCI) signal from CONTROLLER
            cell.count(VPI/VCI) :=0
            counting.cells :=TRUE
        END

25      BEGIN {cell.count timer expires}
            cell.count timer expires
            send cell.count(VPI/VCI) to CONTROLLER
            restart cell.count timer
30      END

```

It will be seen that, as far as the incoming user cell stream is concerned, the activity detector 36 reads the VPI/VCI values in each cell header of the arriving cell stream, and this information is used to update the state table which it maintains for each VPI/VCI value pair. As far as communication with the controller 38 is concerned, the detector 36 informs the controller of a change of state of any VPI/VCI value pair. The controller can inform the activity detector of the timer value *t* to be used. Preferably, the same value of *t* is used for all VPI/VCI value pairs. Cell count information may be obtained from the activity detector 36 by the controller 38, through a "start.cell.count" message. When the activity detector receives this message, it initialises a timer, and counts each cell arrival. At the expiry of each cycle *t*, the activity detector sends the cells count to the controller.

The cells of the user cell stream arriving on input 30 are transmitted without delay to a buffer module 40 where they are stored in first-in, first-out (FIFO) buffer queues, each queue comprising cells having a given VPI/VCI value pair.

5 The detector 36 is non-specific to cell type. Thus, the arrival of any data cells will be detected and can potentially affect the activity state associated with a VPI/VCI value pair, independently of the existence or absence of control or management cells. Buffered cells are

10 fed from the buffer 40 to a shaper multiplexer module 42 prior to being fed to an ATM switch via output 32. Operation of the buffer and shaper/multiplexer modules 40, 42 will be described in more detail below. For the time being, it is sufficient simply to say that the buffer module is capable

15 of signalling to the controller 38 when any buffer queue has reached a predetermined buffer fill threshold. The shaper multiplexer module 42 is responsible for removing cells from the buffer module 40 and transmitting them onwards towards their destination. It includes a multiplexer function and

20 the shaper stores an CR value for each VPI/VCI value pair so that the cell stream fed from the output is shaped to ensure that the capacity of the respective path through the network for each transmission, as determined by the allocated bandwidth, is not exceeded. The controller 38 also controls

25 a feedback module 44 for receiving feedback messages from the network on input 34 and from the controller 38 itself, for onward transmission to the end system 14 via output 35. The functions of the buffer, shaper/multiplexer, and feedback modules 40, 42 and 44 will be described in more

30 detail below. The controller 38 will be considered first.

The purpose of the controller 38 is to signal to the connection admission control (CAC) 18 that an ABR type transmission identified by any given VPI/VCI value pair

35 should have bandwidth in the network allocated or re-negotiated. In this embodiment of the invention, a bandwidth re-negotiation signal is transmitted by the

controller to the CAC 18 such that when the activity detector 36 indicates that a transmission having a particular VPI/VCI value pair is quiet, the bandwidth parameter for that transmission should be set to zero. For
5 a newly active transmission, the bandwidth parameter should be set according to the traffic in the network. This re-negotiation signal is interpreted by the CAC 18 (Figure 1) as a request for available bandwidth, in which case the CAC releases bandwidth to the transmission according to its own
10 control parameters.

The controller 38 signals re-negotiation whenever the activity detector 36 indicates a state change for the transmission in question. The CAC 18 responds by feeding a
15 new CR advice value to the controller 38 which, in turn, updates the shaper/multiplexer module 42 and signals the new CR value to the end system 14 (Figure 1) via output 35. During the period when the CAC is processing a bandwidth re-negotiation or allocation request, the controller 38 sets a
20 default CR in the shaper/multiplexer and feedback modules 42 and 44.

The controller 38 is also responsible for monitoring the actual sending rate of users in cases of transmissions
25 having seized a large effective capacity. This is done by initializing a timer in the activity detector 36 which may have a time-out period of a few seconds, and at the end of this period the activity detector 36 delivers a cell count to the controller 38. Cell counting continues for as long
30 as the signal for the relevant VPI/VCI value pair has a large share of bandwidth in order that the bandwidth is not seized by the signal source, which then transmits at full capacity at first, and subsequently reduces its output to a trickle. The cell count information is processed together
35 with the last CR value advised by the CAC to estimate the network capacity equivalent to the actual sending rate of the user. If this capacity is substantially less than the

currently allocated capacity, it may imply that the user is attempting to keep a large capacity simply by sending a relatively slow trickle of cells to input 30 to keep the activity detector status for the relevant transmission in the active state. Clearly this situation results in inefficient use of the network, and the following steps are performed by the controller 38 when there is a substantial mismatch in the capacity assigned and currently used. Firstly, the shaper 42 is immediately re-set to drop to the actual CR for the relevant VPI/VCI value pair. Next, the feedback module 44 is commanded to issue the actual CR to the user via output 35. Finally, the controller 38 sends a bandwidth re-negotiation signal to the CAC, the signal including a bandwidth value which is set equal to the current equivalent capacity being used by the user. This signal is then interpreted by the CAC as an option to free capacity for other customers. As an alternative, the mismatch may be dealt with by assigning to the shaper and the feedback modules 42, 44 a default CR rather than the actual CR.

As mentioned above, the controller 38 is also arranged to receive a signal from the buffer module 40 when the buffer fill for a given VPI/VCI value pair has reached a given threshold. This signal causes the controller 38 to command the feedback module 44 to issue a so-called resource management (RM) cell, which will be described in more detail hereinafter. The controller 38 may also receive a DBC identity value for each new transmission (identified by a new VPI/VCI value pair) which is established, this DBC identity value being received from the CAC. Alternatively, the DBC may use a default identity if none is supplied.

Pseudocode for the controller appears below:-

35

```
BEGIN
  receive a CR from CAC
  derive minimum.cell.count value
```

21

```

    IF minimum.cell.count  value C>M{WHERE M INDICATES A LARGE
      SHARE OF BANDWIDTH}
      THEN
        send signal to ACTIVITY DETECTOR to start cell.count (VPI/VCI)
5      ELSE
        do nothing
      END

    BEGIN
10      receive cell.count(VPI/VCI) from ACTIVITY DETECTOR
      IF cell.count(VPI/VCI)<minimum.cell.count  THEN
        send actual CR to Shaper
        send actual CR to Feedback
        advise CAC of actual rate
15      ELSE
        do nothing
      END

    BEGIN {Receive a VPI/VCI buffer threshold signal}
20      receive a VPI/VCI buffer-threshold  signal from VPI/VCI-buffer
        advise FEEDBACK to retransmit CR to end system
      END {Receive a VPI/VCI buffer threshold signal}

    BEGIN {Receive a VPI/VCI active signal}
25      receive a VPI/VCI active signal from ACTIVITY DETECTOR
        advise CAC of newly active VPI/VCI
        advise SHAPER of newly active VPI/VCI and default CR
        advise FEEDBACK to transmit default CR in RM cell to end system
      END {Receive a VPI/VCI active signal}
30

    BEGIN {Receive a VPI/VCI inactive signal}
      receive a VPI/VCI inactive signal from ACTIVITY DETECTOR
        advise CAC of inactive VPI/VCI
      END {Receive a VPI/VCI inactive signal}
35

```

Corresponding SDLs for the controller are shown in Figures 6-1 and 6-2. In Figure 6-2, L is typically 1 Mbit/s.

As will be seen from the above pseudocode, the DBC 20 issues
40 a bandwidth re-negotiation request to the CAC for a given transmission (having a given VPI/VCI value pair) when its activity state as indicated by activity detector 36 has changed. This state may be deliberately changing only slowly if the activity detector 36 is set not to notice
45 periods of quiet behaviour which are less than several seconds in length. In response, the DBC 20 receives a signal from the CAC advising the new CR value. The DBC may

also be advised of the correct DBC identity for each transmission.

The controller 38 may be arranged to write a DBC, VPI/VCI
5 identity value pair into the feedback module 44. It is also
arranged to instruct the feedback module 44 to issue a
resource management command for a specific VPI/VCI value
pair. This instruction may also contain an appropriate CR
pair T, τ , (T is an average cell inter-arrival time and τ is
10 a burst tolerance). It should be noted that only one change
in the values specified in an RM cell is sent for each new
VPI/VCI value pair whenever the CAC updates the CR value.
Typically, this may be once every 30 seconds or more in a
public network, and depends upon the sensitivity setting of
15 the activity detector in the DBC 20. It follows that the
required feedback control bandwidth can be relatively small.

As will be seen from the pseudocode, the controller 38
receives signals from the buffer module 40 whenever a buffer
20 fill threshold is reached by cells having a specific VPI/VCI
value pair.

The interface with the activity detector 36 has already been
described.

25

With regard to the controller pseudocode for the interface
with the shaper/multiplexer module 42, the controller 38 is
seen to update the shaper with the current CR value to be
applied to a given transmission whenever it receives CR
30 advice from the CAC or (using the default CR) whenever it
receives advice of a newly active transmission from the
activity detector 36.

The purpose of the feedback module 44 will now be described
35 briefly.

. 23

As mentioned above, the feedback module 44 transmits current CR values (as signalled by the controller 38) to the end system via output 35. The CR is transmitted using a resource management cell as shown in Figure 7. Optionally, one field of this cell is the DBC identity value which is used to enable an end system 14 (see Figure 1) to distinguish between CR advices from different DBCs (e.g. DBCs 20-1 and 20-2 as shown in Figure 2) in the end system to end system path. This DBC identity field is indicated as field 50 in Figure 7. The CR is placed in field 52. This RM cell, like other cells, has a five octet header which contains a PT field 54 indicating that the cell is a resource management (RM) cell.

It is proposed that, if used, DBC identity values are not fixed but are chosen at the time of setting up the transmission path through the network for a given a VPI/VCI value pair. This implies that the CAC 18 assigns a value for the DBC identity for each VPI/VCI value pair, and the feedback module 44 maintains a table of (DBC, VPI/VCI) identity pairs. For example, in Figure 2, public network 10-1 is arranged to choose a DBC identity for a given VPI/VCI pair and signals this information forwards so that public network 10-2 does not chose the same value (e.g. public network 10-1 assigns identity 1, public network 10-2 assigns identity 2, etc). The DBC identity value is stored in a table maintained by the feedback module 44.

The CR field 52 in the RM cell (see Figure 7) contains the CR advice from the CAC which is provided as the average cell inter-arrival time T , plus a burst tolerance τ .

Operation of the feedback module 44 is triggered by the controller 38 (a) when a new CR is advised by the CAC 18, and (b) when the buffer fill level in buffer module 40 corresponding to any VPI/VCI value pair rises above the

buffer fill threshold. A RM cell is then sent to the end system.

The pseudocode for the feedback module 44 is as follows and
5 the corresponding SDL is shown in Figure 8.

```

      BEGIN {Receive a CR}
        receive a CR for a VPI/VCI from Controller
        default_CR:=CR
10      END {Receive an CR}

      BEGIN {RM.cell timer expires}
        RM.cell timer expires
        create RM.cell
15      write default.CR into RM.cell
        send RM.cell to end-system
        restart RM.cell timer
      END {cell arrival from network}

```

20

Next the buffer module 40 will be considered.

The buffer module is shown in more detail in Figure 9. Its purpose is to store incoming data cells on the basis of the
25 VPI/VCI value pairs contained in the cells. Buffering the cells allows an end system 14 (Figure 1) time to respond to a feedback signal from module 44. Another function of the buffer module 40 is to send a signal to the controller 38 when the buffer fill threshold is reached, indicating that
30 an end system is not responding to a feedback signal (this in turn causes the controller 38 to re-send a sustained cell rate signal (CR) to the end system, as mentioned above). The buffer module 40 also drops received cells when the maximum buffer allocation for a given VPI/VCI value pair is
35 exceeded. This is done by buffer overflow.

The size of buffer required for a DBC 20 controlling access to the switching network 10 could be relatively small. For example, if the DBC 20 has a combined input rate from all
40 sources of 150 Mbit/s, then, if the round trip delay to the end system is 100 μ s, there will be less than 35 cells in flight whenever the CR values are changed.

The size of the shared memory area 56 is mainly to cater for changes in the burst tolerance, because a change in this rate leads to only a small number of excess cell arrivals (e.g. around 35 cells). The fixed cell positions assigned to respective VPI/VCI value pairs are designated by the reference numeral 58 in Figure 9. The cells in these positions represent the front cells of a plurality of queues, each queue having its own VPI/VCI value pair. In other words, the queues can be visualised as running horizontally in Figure 9 with the front cells at the right hand side. Cells arriving in the buffer module 40 are placed in the queues in a first-in, first-out (FIFO) order.

Cells are removed from the buffer module 40 when an appropriate signal is received from the shaper section of the shaper/multiplexer module 42, as defined by the buffer module pseudocode which follows:-

```

20      BEGIN {Receive a cell}
        receive a cell
        IF there is room in the buffer THEN
            put cell in buffer
            increment buffer-fill level
        IF buffer-fill level = Threshold THEN
25            transmit buffer-full signal to CONTROLLER
        ELSE
            do nothing
        END {Receive a cell}

30      BEGIN {Receive a fetch}
        receive a VPI/VCI fetch signal from the SHAPER/MUX
        pass cell from buffer to the SHAPER/MUX
        decrement buffer-fill level
        END {Receive a fetch}
35
```

The corresponding SDL appears in Figure 10.

Referring now to Figure 4 in combination with Figure 11, the shaper/multiplexer module 42 operates to remove cells from the buffer module 40 and to transmit them onwards towards their destination via the network switches. Module 42 has two parts which are a multiplexer 60 and a shaper 62. For

each VPI/VCI value pair, the shaper 62 maintains a sustained cell rate (CR) value and a timer.

The cell stream fed to output 32 is shaped by the shaper so that bursts which are not greater than the burst tolerance τ pass without being delayed by the shaper 62. However, the multiplex function may delay a cell if several transmissions represented by different VPI/VCI value pairs are bursting simultaneously. In this case, the multiplexer 60 assigns each active VPI/VCI value pair a fair share of the DBC output bandwidth. It does this by polling active VPI/VCI value pairs in a round-robin fashion. Cells which are waiting for a period equal to or greater than the rate interval T are flagged with a higher priority "cell must go" value. The multiplexer picks up these cells first (see Figure 11). Cells will be forced to wait by the shaper function if bursts arrive which are longer than the burst tolerance credit value. The detailed operation of the shaper/multiplexer module 42 will become apparent from the following pseudocode:-

```

1. BEGIN {STATE = ACTIVE}
   receive a cell.waiting[VPI/VCI] signal from buffer
   IF burst credit ok THEN
25     cell.can.go := TRUE
       STATE := WAIT for multiplexer
   ELSE {burst credit not ok}
       STATE := WAIT for credit timer to expire
   END
30
2. BEGIN {STATE = WAIT for credit timer to expire}
   credit timer expires
   increment burst tolerance credit counter
   cell.can.go := TRUE
35   cell.must.go := TRUE
       STATE := WAIT for multiplexer
   END

3. BEGIN {STATE = WAIT for multiplexer}
40   receive a fetch.cell[VPI/VCI] from multiplexer
   decrement credit counter
   cell.can.go := FALSE
   cell.must.go := FALSE
   STATE := ACTIVE
45   END

```

27

```

4.   BEGIN {STATE =WAIT for multiplexer}
      credit timer expires
      IF credit counter <  $\tau$  THEN
        increment credit counter
5    ELSE
      do nothing
      cell.must.go :=TRUE
      END

10  5.   BEGIN {STATE =SHAPER ACTIVE}
      credit timer expires
      IF credit counter <  $\tau$  THEN
        increment credit counter
      ELSE
15    do nothing
      cell.must.go :=TRUE
      END

20  6.   BEGIN {STATE =SHAPER.CR ACTIVE}
      new CR advised (T,  $\tau$ )
      nextT:= T
      nextcredit:=  $\tau$ 
      END

25  7.   BEGIN {STATE =SHAPER TIMER ACTIVE}
      timer expires
      reset timer (nextT)
      END

30  8.   BEGIN {STATE =MULTIPLEXER ACTIVE}
      output cell timer expires
      index :=pointer
      REPEAT {1st loop of searching for cell.must.go}
        increment index
35    IF cell.must.go[index] THEN
      pointer:= index
      fetch cell[index] from buffer
      send fetch cell signal to SHAPER
      STATE :=MUX.ACTIVE
40    ELSE
      IF index =max.buffer.size THEN
        index :=0
      UNTIL index =pointer

45    REPEAT {2nd loop of searching for cell.can.go}
      increment index
      IF cell.can.go[index] THEN
        pointer:= index
        fetch cell[index] from buffer
50    send fetch cell signal to SHAPER
      STATE :=MUX.ACTIVE
      ELSE
      IF index =max.buffer.size THEN
        index :=0
55    UNTIL index =pointer
      send no.cell.waiting.signal

```

END

It will be understood that when the DBC 20 requests a change in the bandwidth allocated to a particular transmission, the
5 CAC must control other traffic in the network so that the network capacity is used most effectively. The description which follows deals with connection admission control methods for overcoming the problem of traffic rebalancing.

10 Two connection admission control strategies will now be described. Both tackle the problem of rebalancing traffic. In other words, when a transmission becomes quiet or newly active, it is necessary to determine how many other control
15 messages need to be generated for other transmissions. The object is to make this number of control messages as small as possible.

The first strategy involves a relatively simple connection admission control method which involves no actual
20 rebalancing. In this method, a newly active transmission (VPI/VCI value pair) is given a single sustained cell rate (CR) which is retained until the transmission goes quiet again. Only when it is subsequently reactivated will the transmission get a different CR. This means that a quiet
25 signal relating to one VPI/VCI value pair will cause no control signals to be generated for other VPI/VCI value pairs which were sharing capacity with it.

This is combined with a filling method which involves (i)
30 giving a first newly active connection an effective capacity which is half of the total available capacity; (ii) giving the next newly active connection an effective capacity which is half of the remaining capacity; (iii) giving the next newly active connection an effective capacity which is half
35 of the still remaining capacity; and so on. This method is applied link-by-link over the entire route identified by the VPI/VCI value pair, and whichever yields the lowest

effective capacity is the determinant of the CR fed back to the DBC 20.

It follows that a newly active signal having one VPI/VCI value pair generates no control signals for the other VPI/VCI value pairs which are sharing the capacity.

Since the DBC 20 is designed such that a user can only maintain a large effective capacity on the network so long as the VPI/VCI value pair remains in the active state in the activity detector 36 (Figure 4), and the cell rate generated by the customer is close to the effective bandwidth value (refer to the cell-counting function of the activity detector described above), it follows that users can only hold onto large effective bandwidths for as long as they are prepared to be charged for the proportionally larger loads which they are submitting.

This method is fair to users in the sense that, over a sufficiently long period, no user is systematically given a poorer capacity.

However, it is desirable in some circumstances to increase the number of users who are able to secure relatively large bandwidth allocations and this can be catered for by a second, modified method as follows.

In this case the underlying principle is that, if an active signal causes control signals for other VPI/VCI value pairs, let the signal be limited to only one per link, namely the richest (largest capacity) VPI/VCI value pair. This can be described as a limited rebalancing method or a "take-only-from-the-richest" (Robin Hood) method.

This can best be illustrated with an example filling method:-

- (i) the first newly active VPI/VCI value pair is assigned an effective capacity equal to half of the total available capacity;
- 5 (ii) the next newly active connection is assigned half of the remaining capacity plus a fifth of the effective capacity of the first VPI/VCI value pair (i.e. the current richest);
- 10 (iii) the next newly active connection is assigned half of the remaining capacity plus one fifth from the current richest; and so on.

To illustrate this process, it may be imagined that there is
15 a single link with a capacity of 100 Mbit/s. The above steps then result in the following exemplary steps:-

- (i) the first newly active VPI/VCI value pair gets 50 Mbit/s and there is 50 Mbit/s remaining;
- 20 (ii) the next VPI/VCI value pair gets half of the remainder (which yields 25 Mbit/s) plus a fifth from the first, which means that the first now has 40 Mbit/s, and the second has 35 Mbit/s;
- 25 (iii) the next VPI/VCI value pair gets half of the remainder, which yields 12.5 Mbit/s plus a fifth from the first, so that the first now has 32 Mbit/s, the second still has 35 Mbit/s, the third has 20.5 Mbit/s, and so on.
- 30

Note that more of the users are now getting large capacities, but there is only one extra control message to send on the link. There is thus a limited rebalancing or
35 "Robin Hood" strategy.

To extend the method to a route with many links, the above process is repeated link-by-link. Whichever link yields the lowest effective capacity is the determinant of the CR value sent back to the dynamic bandwidth controller (DBC). Now, 5 using this value of effective capacity, the CAC assigns it link-by-link by taking half of the remaining capacity on that link, and any extra which is needed is taken from the richest VPI/VCI value pair on that link. Consequently, this generates at most one additional CR control message per link 10 for each VPI/VCI active signal sent to the network. A quiet signal still generates no additional control messages.

This strategy also makes it impossible for a user to hold onto a very large capacity when others become active. In 15 addition, as many users as possible are given a reasonably large capacity while keeping the complexity of traffic rebalancing to a minimum.

A second embodiment of the invention will now be described 20 with reference to Figures 13 to 15. In this case, operation of a dynamic bandwidth controller (DBC) in combination with a usage parameter controller (UPC) 22 is described, using a modified DBC configuration.

25 It will be recalled that each ATM cell has an information field of 48 octets and a header of 5 octets.

The cell header may include a cell loss priority bit, allowing each cell to be identified as having high priority 30 or low priority. When congestion occurs, that is to say when elements within the overall switching network reach capacity, the network is arranged such that low priority cells are discarded in preference to high priority cells. Furthermore, in any operational network, the network itself 35 would be designed such that, under normal operating circumstances, high priority cells would always be allowed through the network, so that congestion, resulting in the

loss of low priority cells, would only occur during periods of peak demand. For some forms of data transmission, such as 64 kbit/s voice, cell loss may be acceptable. However, in other circumstances, it may be highly undesirable for
5 cells to be lost, particularly when non-redundant data is being transmitted, under which circumstances a user would expect traffic controls to ensure that all data cells were transmitted through the network with very low losses.

10 The arrangement of Figure 13 forms part of a network which may operate in several ways, although two preferred modes of operation will be identified. One mode is to provide permanent connections between users, of the leased-line type. In this mode of operation, a communications channel
15 between a first user and a second user is established, and thereafter this channel remains permanently connected. In a second mode of operation, the network operates in a switched-like way, in that signalling commands are issued by a user. These signalling commands, are in turn, interpreted
20 by the switch 12, resulting in the connection being established. However, in some data networking applications, it is envisaged that, once the connection has been established, it would remain in place for a significant period, possibly days. Therefore, such a link differs
25 significantly from connections made via the public switch telephone network (PSTN), for example, where connection time is usually measured in minutes. In this second mode, then, the network provides a service similar to that provided by leased lines. In its first mode of operation, which
30 provides the basis for the embodiment now to be described the end system behaves as if leased line were provided between communicating sites. Under this preferred mode of operation, it is unnecessary for the end system to include any mechanisms for generating signalling, allowing the first
35 user immediately to initiate communication with the second user.

Thus, signals generated by the first user will be supplied to the switching environment from the usage parameter control device 22 and will be directed to the second user via the DBC 20 and several switches such as switch 12. At
5 the location of the second user, user equipment is provided for the transmission and reception of traffic. Therefore, transmission between users is only in the form of data, with other communication channels being required for the first user to effect communication of voice and video signals via
10 the network.

The broadband switch 12 has a plurality of input ports and a plurality of output ports. Thus, output signals from a user are supplied to an input port 70 of the broadband
15 switch 12, and the switch is arranged to supply cells received at that port to an output port 72, which in turn directs the cells via similar switching devices to the second user. Similarly, output cells from the second user are ultimately supplied to the broadband switch 12 at an
20 input port 74, allowing the switch 12 to direct the cells to an output port 76.

The establishment of logical paths within the broadband switch 12 is effected by a network management computer 78,
25 which is also responsible for overseeing the operation of a plurality or similar broadband switches.

Communication between the network management computer 78 and the broadband switch 12 is made via a respective element
30 manager 80, which is switch-specific and provides an interface between the switch and the network management computer. Thus proprietary broadband switches from a plurality of manufacturers may be configured within an overall network interfaced via respective element managers
35 80.

Traffic control of the broadband switch 12 is effected via a call controller 82. Thus, it is necessary for the call controller 82 to receive information defining the bandwidth requirement for any particular logical connection, thereby
5 ensuring that suitable operation of the broadband switch 12 is made under the control of the call controller, so as to minimise cell loss. The arrangement should be configured such that cells identified as being of a high priority are never lost, and low priority cells are only lost when the
10 broadband switch 12 is overloaded when congested.

The network includes usage parameter control devices 22, which conform to the operations defined within Recommendation I 371. Thus, a usage parameter control
15 device 22 is provided for each user, and is programmed in accordance with the required level of service required by the user. Thus, a user's usage contract will specify a first traffic threshold, below which, all cells identified as being a high priority will remain at this priority,
20 thereby ensuring their successful transmission through the network. However, once traffic exceeds the first threshold, some cells will have their priority downgraded to low priority, such that it cannot be guaranteed that transmission will successfully take place through the
25 network. Thus, a user will have a guaranteed bandwidth, beyond which traffic may be accepted, but transmission through the network cannot be guaranteed.

The usage parameter control device 22 defines a second
30 traffic threshold, beyond which cells will be removed from the data stream altogether, thereby ensuring that the bandwidth requirement made of the switch never exceeds, for that particular user, the traffic level defined by the second threshold.

35

Thus, in normal operation, a user would aim to stay below the first threshold, thereby ensuring that all cells are

transmitted through the network, However, if (possibly through some unexpected occurrence) traffic demand increases, the network may have capacity to accommodate this increase in traffic, thereby maintaining data integrity for the customer, while at the same time being provided with a mechanism for removing offending cells, thereby ensuring that cells to which bandwidth has been guaranteed are transmitted through the network.

10 Thus, within each broadband switch 12 and under the control of the call controller 82, low priority cells may be rejected when congestion occurs. However, as previously stated, the overall network would be configured such that high priority cells would be guaranteed transmission through
15 the network.

In the first preferred embodiment of the invention, a user may remain permanently physically connected to the network, and may be under the impression that the network is
20 permanently configured so as to make bandwidth available, even when the user does not have a requirement to make use of this available bandwidth. However, the user is only charged when an actual use of the network is made, preferably by charging the user on the basis of the number
25 of cells transmitted over the logical link. Thus, prior to transmission being invoked by the user, the broadband switch 12 may be actually allocating no bandwidth whatsoever for transmission between the input port and the output port. Furthermore, it is not necessary for any special activity to
30 be made by the user in order for bandwidth to be allocated. A level of bandwidth will be assigned to the user when required.

After transmitted cells have passed the usage parameter
35 control device 22, they are directed to a dynamic bandwidth controller 22. It may take a short interval for the call controller 82 to find the required level of bandwidth

required by the bandwidth controller 20 so that, during the interval required by call control to process the bandwidth request, all transmitted cells are downgraded to low priority.

5

In addition to the user generating signals, which, as previously stated, are supplied to the input port 70, signals are also transmitted back to the user via the output port 76. These output signals are also supplied to the dynamic bandwidth controller 20, allowing said controller to transmit signals back to the user. Thus, in particular, the dynamic bandwidth controller 20 may issue a signal back to the user, instructing the user to reduce its data output to avoid the risk of cells being lost.

15

The dynamic bandwidth controller 20 is shown in more detail in Figure 14, and includes an activity detector or cell monitor 36, a controller or processor 38, a feedback module 44 and an attribute data store 84.

20

Data is transmitted to the port 70 and similar data is received from the port 76 in the form of discrete cells, having forty-eight octets of user information plus a five octet header.

25

As mentioned above the last 8-bit octet of the header is a header error check field which provides a degree of redundancy, by which error checking may be performed on the header information. Thus, the principal reason for providing the header error check field is to ensure that the header information is correct.

Within the cell monitor 36, the header error check field is also used to identify the start of a cell.

35

In addition to customer originating data being transmitted in the form of cells, it is also possible for the processor

38 to communicate with customer equipment and sub-switching networks within an overall switching environment by generating control information cells.

5 As previously stated, the usage parameter controller 22, shown in Figure 13, is capable of modifying the priority of transmitted cells. The traffic thresholds of the usage parameter controller 22 may be adjusted in response to signals received from the network manager 78.

10

In previous systems, signals would only be supplied to the usage parameter controller 22 when a customer's contract was modified, reflecting a change in the customer's requirement for bandwidth. However, in the present system, the
15 allocation of bandwidth to customers is dynamically controlled so that, in its non-active state, a transmission path having zero bandwidth capacity is effectively allocated for a customer. Thus, in order to reflect this state, an instruction is issued by the network manager 78 to the usage
20 parameter controller 22 setting the first threshold level to zero. Under these conditions, all cells received by the usage parameter controller will be modified to cells of low priority, so that it cannot be guaranteed that the cells will be transmitted through the network. However, during
25 these non-active periods, it is assumed that no demand is being made of the network and that switching capacity may be allocated to other users.

In the preferred embodiment, it is not necessary for a
30 customer to perform any signalling functions in order to initiate a call. The data itself is identified by the dynamic bandwidth controller 22, and triggers the generation of bandwidth requests to the network management or call control.

35

Referring to Figure 14, low priority cells, during the initiation of transmission, are supplied to the cell monitor

38

36 which writes these cells to a first-in-first-out (FIFO) shift register for re-transmission, with minimal delay, to the port 70. The shift register includes outputs which supply values to combinational logic which, as previously described, is capable of identifying cell phase, with reference to the header error check field.

On identifying the presence of cells within its shift register, the cell monitor 36 generates a transmission active signal, which is supplied to the processor 38 along with an indication of the VCI and the VPI associated with the cells being transmitted. In response to this information, the processor 38 will interrogate the attribute data store 84 so as to identify the bandwidth allocation associated with said identifiers. In response to this enquiry, data is supplied back to the processor 38, identifying the bandwidth provision for the particular communication identified by the call monitor 36.

The processor 38 is arranged to configure control cells for transmission to other elements in the network, and in response to the "transmission active" signal generated by the cell monitor 36 and the data read from the attribute data store 84, a message is assembled for transmission to the network manager 78, requesting the establishment of bandwidth within the broadband switch 12, facilitating communication between the input port 70 and the output port 72. In response to this request, the network manager 78 will determine whether this amount of bandwidth is available and, if so, bandwidth is assigned to the connection.

In addition, the network manager 78 also transmits a message to the port 76, for the usage parameter controller 22 which, in response to this message, adjusts its threshold values such that cells are no longer marked as low priority provided that the rate is less than, or equal to, the bandwidth now granted to the connection.

In some situations, a customer may wish to transmit important data where a guarantee is required that the data will be received at its destination. If the terminal transmits a sequence of trial cells all marked as high
5 priority then, under the scenario described above, it can be appreciated that initial cells will have their priority level downgraded by the usage parameter controller 22 until a state has been reached such that bandwidth allocation is guaranteed through the broadband switch 12. Thus, at the
10 receiver, a succession of low priority trial cells may be received prior to the reception of high priority cells, which are only to the reception of high priority cells, which are only transmitted after bandwidth has been granted.

15 Thus, where a customer requires a guaranteed communication to be established, it is possible for a receiver to examine the status of received cells, and then to issue a signal back to the originating customer when cells of high priority are being received. The receipt of high priority cells will
20 indicate that high priority cells will be maintained and, under this condition, the originating customer may then transmit the high priority data, secure in the knowledge that the priority cells will remain as such throughout the transmission through the network; it only being the initial
25 start up period where cell priority downgrading occurs. Thus, the network automatically responds to a request for bandwidth provision without undergoing an initiation or signalling routine. The trade-off for this functionality is that the priority level of initial cells will be downgraded,
30 and the duration over which this downgrading occurs will depend upon the capacity of the network to establish the bandwidth requirement and thereby issue modifying signal to the relevant usage parameter control devices; therefore, bandwidth provision is guaranteed.

35

The cell monitor 36 is also capable of detecting the absence of cells passing therethrough. As previously described,

provision is made for the identification of cell transmission, and the cell monitor 36 includes a time-out circuit which will ensure that a periodic investigation is made as to whether cells are continuing to be transmitted.

5 When transmission ceases, the cell monitor 36 will effectively time-out and supply a signal to the processor 38 to the effect that transmission has ceased. In response to this signal, the processor 38 will configure a message and transmit this cell to the network manager 78 via the cell

10 monitor 36. In response to this message, the network manager 78 will allow re-allocation of the bandwidth within the broadband switch 12, and issue a message to the usage parameter controller 22, again re-setting the threshold level to zero, such that, on the next iteration, initial

15 cells will have their priority level reset to low priority.

It will be appreciated that many logical connections may be achieved via common physical links. Thus, at a particular input to the overall network, an input port, such as the

20 port 70, may receive cells relating to a plurality of logical connections; and, similarly, the dynamic bandwidth controller 20 will control bandwidth allocation for each of these logical connections.

25 The allocation of bandwidth has been described for the broadband switch 12 shown in Figure 13, which is the first broadband switch encountered on entry to the overall network of the type shown in Figure 1. The bandwidth request made by the dynamic bandwidth controller 20 will also result in

30 interrogations being made of each switch through which the connection passes in order to establish the availability of bandwidth. Thus, provision may have to be made at a plurality of broadband switches, before a message is returned to the dynamic bandwidth controller 20, confirming

35 that bandwidth has been allocated, thereby ensuring that high priority cells are transmitted through the overall network.

As previously stated, the header information is identified by the processor 38 and, in response to this information, attribute data is read from the attribute data store 84, identifying the bandwidth allocated to that particular channel. Thus, for example, the attribute data store 84 may identify the channel as a 1 Mbit/s channel, a 5 Mbit/s channel or a 10 Mbit/s channel etc, depending upon the level of service contracted to the customer.

- 10 A further feature provided by the dynamic bandwidth controller 20 is that of allowing a customer access to whatever bandwidth is available at the particular time when a request is made. Thus, when cells arrive from a customer to whom this provision has been made available, an indication is supplied to the processor 38, from the attribute data store 84, to the effect that the cells are to be given whatever bandwidth is available for providing a channel between the two communicating stations.
- 20 The processor 38 will construct a message and transmit this cell, identifying the customer's request to the network manager 78. In response to this request, the network manager 78 is required to determine the level of available bandwidth from the communicating source to the communication destination. The available bandwidth over the logical channel will be restricted by whichever physical link has the minimum available bandwidth. Thus, if a communications channel is constructed over three serial physical links, a first of which has 10 Mbit/s available, a second of which has 10 Mbit/s available while a third only has 2 Mbit/s available, the bandwidth available over the combination is only 2 Mbit/s and the remaining 8 Mbit/s at the first and second stages of the link cannot be employed because a bottleneck exists at the station which only has 2 Mbit/s available. Thus, in response to the request made by the processor 38, the network manager 78 will determine the

total bandwidth available, and return this information back to the dynamic bandwidth controller 20.

A complication arises when provision is made for allocating
5 whatever bandwidth is available in response to a customer request. At the time the request is made, the requesting customer is not aware what actual level of bandwidth will be available. Furthermore, the network is not aware of the extent to which the requesting customer will require
10 bandwidth. Thus, the network will provide whatever bandwidth is available; and, under some circumstances, this may be insufficient for the customer's requirements. Under these conditions, it is necessary for the network to issue instructions back to the customer, in the form of suitably
15 constructed cells, so as to inform the customer that congestion is imminent, and that action must be taken to reduce the level of traffic supplied to the network.

Cells of this type, instructing a transmitting customer's
20 terminal to reduce its output traffic, are generated by the processor 38 and supplied to the originating terminal via the feedback controller 44. In order to determine whether such a feedback instruction requires to be generated, the cell monitor 36 includes buffers arranged to buffer
25 incoming data of this type, and to generate overflow signals when it is detected that data is being written to the buffers at a faster rate than it is being read from the buffers.

30 The cell monitor 36 of Figure 14 is shown in more detail in Figure 15. It should be appreciated that, by its very nature of being capable of transmitting cells in accordance with the available bandwidth, cells of this type will be transmitted at different rates, depending on the bandwidth
35 available when the particular request is made, and that separate buffering devices be provided for each particular rate.

In theory, the level of available bandwidth would be continuously variable and, in response to such a request, any number (within the definition of the system) could be returned to the requesting processor. Under such
5 circumstances, in order to make full advantage of the available bandwidth, it would be necessary to provide buffering devices for every possible transmission rate, or to provide variable rate buffers for each logical channel being transmitted through the system. Each of these
10 alternatives is undesirable.

As shown in Figure 15, four physical buffers 86, 88, 90 and 92 are provided. The first buffer 86 is arranged to buffer cells being transmitted at 1 Mbit/s, the second buffer 88 is
15 arranged to buffer cells being transmitted at 2 Mbit/s, the third buffer 90 is arranged to buffer cells transmitted at 5 Mbit/s, and the fourth buffer 86 is arranged to buffer cells transmitted at 10 Mbit/s. In practice, the actual number of buffers provided will depend upon operating
20 requirements, as well the actual data rates for which they accommodate cells.

Cells are transmitted from customer equipment and supplied to the port 70 as a serial stream. However, in the cell
25 monitor 36, the serial stream is converted into parallel bit octets by a serial-to-parallel converter 94, which in turn supplies these octets to a parallel shift register 96. The shift register 96 is tapped, allowing a plurality of octets to be read simultaneously by a logic circuit 98 which, as
30 previously described, is arranged to identify the presence of data transmission, the position of headers within cells, and to supply header information to the processor 38. The processor 38 will, in turn, identify the stored bandwidth allocation for the particular cell being transmitted and, in
35 response to this determination, will supply control signals to a switch 100.

If a cell is identified as belonging to a logical communications channel for which the required bandwidth has been specified, the switch 100 is arranged, in response to signals received from the processor 38, to supply cells
5 directly to a multiplexer 102. However, if the cells detected are of the type which will result in a request for "whatever bandwidth is available" to be made available for their transmission, cells supplied to the switch 100 will be directed to one of the buffers 86 to 92, rather than being
10 supplied directly to the multiplexer 102.

After a request has been made by the processor 38 to provide "Whatever bandwidth is available", the network manager 78 will supply information back, in the form of an
15 appropriately addressed message, to the processor 38 identifying the level of bandwidth which is available. When a level of available bandwidth is received by the processor 38, it will not immediately make this bandwidth available for transmission. In preference to making whatever
20 bandwidth is available actually available for transmission purposes, the processor 38 reduces the level of bandwidth available to a level equivalent to the nearest lower value for which a buffer has been provided. Thus, referring to the buffer values provided in Figure 15, if the available
25 bandwidth comes back as being greater than 10 Mbit/s, the actual bandwidth made available is 10 Mbit/s, and the buffer 92 is used. Similarly, if the available bandwidth is greater than 5 Mbit/s, but less than 10 Mbit/s, 5 Mbit/s is made available and the buffer 90 is used. Similarly, if the
30 actual available bandwidth is between 2 Mbit/s and 5 Mbit/s, the actual allocated bandwidth provision is set to 2 Mbit/s and the buffer 88 is used. Similarly, if the actual bandwidth available is between 1 Mbit/s and 2 Mbit/s, the actual bandwidth allocated is set at 1 Mbit/s, allowing the
35 buffer 86 to be used. Finally, if the actual available bandwidth is less than 1 Mbit/s, it is assumed that no

bandwidth is actually available and no provision will be made for facilitating high priority data transmission.

Thus, as described above, the processor 38 will convert the
5 actual level of bandwidth available to one of the processable transmission rates and supply a suitable signal to the switch 100 accordingly. For example, if it is determined that the actual available bandwidth is 3 Mbit/s, a signal identifying this level of provision is supplied to
10 the processor 38 from the network manager 78. In response to this signal, the processor 32 selects one of the processable transmission rates which, in this example, will be 2 Mbit/s, thereby allowing transmission to occur at high priority up to a limit of 2 Mbit/s. A signal is supplied to
15 the switch 100 from the processor 38, thereby directing cells from this link to the buffer 88. The cells are clocked through the 2 Mbit/s buffer 88 in a first-in-first-out fashion, ultimately resulting in said cells being supplied to the multiplexer 102. It should be appreciated
20 that the buffer 88 will also receive 2 Mbit/s cells from other logical links, all of which will be clocked out of that buffer at the appropriate rate. At the multiplexer 102, these cells are combined with cells received from the other buffers and cells received directly from the switch
25 100 for transmission, during appropriate time slots, to the input port 70. As shown in Figure 15, the multiplexer 102 is also arranged to receive cells directly from the processor 38 which, as previously described, are required for housekeeping purposes within the switching environment.

30

Now, to summarise the application of the embodiment described above with reference to Figures 13 to 15 within existing networks, a terminal is connected, via a physical link, to a local exchange, which includes interconnection
35 function, control functions for cell processing, management functions for managing resources and usage parameter control functions for policing traffic entering the public network.

All of these functions are common to existing ATM based machines. However, it is proposed that the local exchange includes an additional function which is termed the dynamic bandwidth controller (DBC), which can interact with the management and/or call control functions.

The DBC functions may be summarised as follows:

- I. To monitor specific private virtual circuits (PVCs) within the link which connect the terminal to other terminals via interconnection functions and to detect whether traffic activity commences on any PVC.
- II. An option of the cell monitor 36 is to set the CLP bits to low priority on all cells of a monitored PVC whenever new activity has been detected and while no bandwidth has been assigned to the PVC other than the default bandwidth. Alternatively, to conform to standards which dictate that the CLP bit shall only be changed by the usage parameter control (UPC), so as to transmit a control message to network management/call control once inactivity has been detected on a PVC, requesting that the UPC parameter is updated, thereby ensuring that all subsequent cells are passed (up to an agreed maximum rate) but are marked as low priority, "violation tagged".
- III. When traffic activity is detected on a PVC, to generate signalling messages to the call control function and/or management functions, and to request that the connection attributes are changed, that is, bandwidth is assigned to the PVC which is either a fixed value agreed with the customer for that PVC, or is the highest available bit rate.
- IV. To receive signalling messages from the cell control function and/or the management function,

acknowledging the "change attribute" request and indicating that bandwidth had been granted for this PVC, if any.

- 5 V. To stop overwriting of the CLP with a low priority indication whenever bandwidth has been granted. Alternatively, to conform with standards which state that the CLP bit shall only be modified by the UPC to generate a control message to network management/call control whenever bandwidth has been granted for a PVC, requesting that the UPC parameters are updated, so that cells are passed up to given maximum rate with no modification, ie violation tagging, of the CLP.
- 10
- 15 VI. To provide feedback control to the terminal, as necessary, to keep the traffic associated with this PVC at a rate which is not greater than can be supported by the granted bandwidth.
- 20
- VII. To detect when traffic activity ceases for a period of more than n milliseconds and, if this occurs, to generate a signalling message informing the call control function and/or the management function that the connection attributes should be changed, such that zero bandwidth is assigned.
- 25

The cell monitor function decides, for each PVC, whether traffic activity has commenced or ceased. Cells transmitted by a terminal, or other PVC end system, are passed to the DBC via an interface. The cell monitor function may include serial-to-parallel conversion of the arriving digital signals, and a shift register to store all or part of the cell header, to facilitate processing of the VPI/VCI and the CLP bit. The cell monitor function includes the capability that it can set the CLP bit to low priority on any cell. However, this can be turned off by the processor function.

30

35

Prior to passing the cell onwards towards the network, via an interface, the cell monitor function can perform parallel-to-serial conversion of the digital signals, as required.

5

The cell monitor function informs the processor function whenever traffic activity has changed on a PVC. The processor function generates appropriate control signals to activate/deactivate the CLP over-writing function within the cell monitor. The CLP over-writing function is activated on a particular PVC whenever call control or network management informs the processor that no bandwidth is assigned on that PVC. Alternatively, the processor function generates a control signal which is transmitted to request the call control or network management to update the UPC parameter, so that all cells are violation tagged on this PVC. The CLP over-writing function is deactivated on a particular PVC, whenever call control or network management informs the processor that guaranteed bandwidth has been assigned to the PVC. If required, the processor function will also generate a request to call control or network management to the effect that UPC parameters have been changed to stop violation tagging.

25 The processor function also generates control signals to call control and network management, which request bandwidth changes to the PVC. Bandwidth request and acknowledge signals are sent and received via an interface. The correct generation of these signals by the processor relies on data stored in the PVC attributes data store, which contains information on the amount of guaranteed bandwidth required or whether any available bit rate can be assigned up to a given maximum value and above a given minimum value.

35 Cells which travel in the opposite direction towards the terminal are passed to a DBC feedback function via an interface. The purpose of this function is to instruct

terminals to stop or transmit at any speed which does not infringe the peak rate or sustained cell rate specified in traffic contract monitored by the UPC. The feedback function may pass these signals back to the terminal either
5 by modifying the genetic flow control (GFC) field of any cell header or by inserting a control cell in place of an unassigned cell, when this arrives.

The feedback function may include serial-to-parallel
10 conversion of the arriving digital signals, and a shift register is provided to allow the processing of cell headers and parallel-to-serial conversion of output digital signals.

The processor function informs the feedback function of the
15 correct rate at which to apply feedback signals to the terminals. No feedback signals are required to be generated when the customer has requested that a fixed bandwidth is to be assigned, rather than any available bit rate. In this case, it is assumed that the terminal always transmits at
20 the fixed rate, which is monitored by the UPC. Therefore, no further control function is required to keep the terminal transmitting at the correct rate.

However, for customers who have requested that any available
25 bit rate should be assigned to the PVC, between given minimum and maximum values, it is assumed that the UPC only checks that the maximum has not been exceeded.

Cells from available bit rate PVCs, when arriving at the
30 DBC, are routed internally to one of several buffers within the cell monitor function. The cells from fixed bit rate PVCs need not be routed by any buffer, since no additional bit rate checks are necessary, other than as provided by the UPC, in which case there are two distinct delay
35 characteristics. Firstly, fixed bit rate PVCs carrying bursty traffic with low delay, suitable for variable bit rate video and, secondly, available bit rate PVCs carrying

bursty traffic with a large buffer delay, suitable for data which is not delay-sensitive.

The buffers within the cell monitor function are each
5 associated with a traffic class, based upon an output bit rate. All available bit rate PVCs are assigned to one of these classes, based on the bandwidth granted by the call control and network management.

10 When requesting any available bit rate for a PVC, the DBC may either supply the acceptable class rates as parameters of the request, or start with a request for the highest class rate acceptable for the PVC, as stored in the PVC attribute data store and, if call control or network
15 management does not grant that request, a second request is generated for the lowest class rate acceptable for the PVC. If this request is granted, no further requests are made, and the PVC is assigned to that class. Until an available bit rate PVC is assigned to a class, no cells are diverted
20 to any buffer, and all cells are marked with low priority. Furthermore, cells arriving at the DBC which are not diverted to a buffer, either from fixed bit rate PVCs or from available bit rate PVCs, which have not yet been granted any bandwidth, have priority over cells in a
25 transmit buffer. A transmit buffer is a common buffer to which cells are transferred from class buffers, prior to output.

This rule of transmission priority is necessary to ensure
30 cell sequence integrity of available bit rate PVCs and to ensure that fixed bit rate PVCs have minimum delay. To understand the cell sequence integrity guarantee for available bit rate PVCs, it is clear that, if the first R cells are transmitted prior to bandwidth being guaranteed,
35 and if cell R+1 and all remaining cells (until inactivity is detected) are diverted to a class buffer then, because of

the priority transmission rule above, cell R+1 cannot be transmitted before the Rth cells.

If n PVCs are in the 2 Mbit/s class, cells are transmitted to the transmit buffer at a rate of n times 2 Mbit/s. This ensures that available bit rate PVCs conform to the rate expected by the network. If the terminal is transmitting faster than this rate, the effect will be that the class buffer begins to fill, but there will be additional unexpected traffic entering the network beyond the DBC.

The cell monitor functions provides an indication to the processor function whenever a class buffer has been filled above a pre-assigned threshold, or when it has been filled below another pre-determined threshold. When the processor function receives an indication that a class buffer is filled above said first threshold, it instructs the feedback function to deliver "stop transmitting" signals to the appropriate terminals. The PVCs identities of the appropriate terminals are passed to the feedback function from the processor function, and the feedback function may then insert a series of control cells with the same PVC values which instruct all those terminals to stop. Alternatively, it may stop all PVCs using GFC, if there is a large number of PVCs to stop, for example, if there is only a small number of different classes.

When the processor function receives an indication that a class buffer is filled below the second threshold, it instructs the feedback function that the cell transmission can start again on those PVCs. In this case, the feedback function will no longer transmit control cells to those PVCs or said function will change the GFC setting if it was controlling a large number of PVCs.

CLAIMS

1. A broadband switching system having at least one
5 ingress for connection to a respective signal source
and at least one egress for connection to a selected
signal receiving system, the switching system having
at least one switch for transmitting information-
10 carrying asynchronously transferred data cells from
the ingress to the egress, system control means for
accepting and establishing a connection between the
ingress and the egress via the said switch, and
bandwidth control means arranged to detect the said
15 information-carrying data cells received at the
ingress and, automatically in response to such cell
detection, to cause the system control means to
allocate bandwidth for the transmission of the cells
to the egress.
2. A system according to claim 1, characterised in that
20 the system control means includes means for
identifying available bandwidth for the said
allocation of bandwidth.
3. A system according to claim 2, characterised in that
25 the bandwidth control means includes means operable
to receive from the identifying means a control
signal representing the available bandwidth.
4. A system according to claim 2 or claim 3,
30 characterised in that the system control means
includes a congestion sensor for transmitting a
congestion signal to the bandwidth control means, and
in that the bandwidth control means includes feedback
means for transmitting a reduce-traffic-level
35 instruction to the ingress for the signal source.

5. A system according to any of claims 1 to 4, characterised in that the bandwidth control means include a transmission limiter having storage means for storing a maximum output cell rate value derived from a control signal received by the limiter from the system control means, and means for limiting the transmission of cells to the said switch when the rate at which cells are received at the ingress exceeds the maximum output cell rate value.
6. A system according to claim 5, characterised in that the limiter includes a buffer for delaying transmission of cells to the switch.
7. A system according to claim 5 or claim 6, characterised in that the bandwidth control means include feedback means arranged to transmit the maximum output cell rate value to the ingress for transmission to a signal source generating the data cells.
8. A system according to any of claims 5 to 7, characterised in that the ingress is arranged to receive data cells from a plurality of signal sources and in that the transmission limiter is operable to transmit cells from one of the sources at a rate which is dependent on the number of cells from one or more of the other sources stored by the device.
9. A system according to any preceding claim, characterised in that the ingress is arranged to receive data cells from a plurality of signal sources, in that the bandwidth control means include a cell monitoring stage operable to read routing identifiers associated with the cells and to update activity status values associated with the routing identifiers, and in that the transition of an

activity status value relating to a said identifier from an inactive to an active status causes the system control means to allocate bandwidth for cells associated with the said identifier.

5

10. A system according to any preceding claim, characterised in that the bandwidth control means are arranged to detect the rate at which cells are supplied to the ingress.

10

11. A system according to claim 10, characterised in that the bandwidth control means are arranged such that if the allocated bandwidth is more than that required to transmit the cells at the monitored rate, the system control means are caused to allocate less bandwidth to the cells.

15

12. A system according to claim 1 or claim 10, characterised in that the bandwidth control means are arranged to read a bandwidth-identifying portion of the incoming cells and to cause the system control means to allocate bandwidth to the cells according to the bandwidth identified by the said portion.

20

- 25 13. A system according to any preceding claim, characterised in that the bandwidth control means are arranged to limit the rate at which the incoming cells are accepted onto the system according to a default bandwidth until allocation of bandwidth to the said incoming cells by the system control means has occurred.

30

14. A system according to claim 13, characterised in that the bandwidth control means include means for assigning a low priority status to the said incoming cells while no bandwidth has been allocated to the cells other than the said default bandwidth.

35

15. A system according to any preceding claim, characterised in that the bandwidth control means include a cell counter for counting cells received from respective signal sources coupled to the bandwidth control means thereby to generate charging signals for customer billing.
16. A system according to claim 6 or claim 7, characterised in that the buffer has means for detecting filling of the buffer to a predetermined threshold level, and in that the feedback means are responsive to the buffer to cause a re-transmission of the maximum output cell rate value for the associated signal source when connected to the ingress when the buffer has been filled to the threshold level.
17. The use in a broadband switching system for the transmission of information-carrying asynchronously transferred data cells from an ingress of the system to an egress of the system, of bandwidth control means arranged to detect incoming cells received at the ingress and, automatically in response to such cell detection, to cause system control means of the system to allocate bandwidth for the transmission of the cells to the egress.
18. A method of operating a broadband switching system for transmitting information-carrying asynchronously transferred data cells from an ingress of the system to an egress of the system via at least one switch, the method comprising controlling the bandwidth for transmitting the cells by detecting the cells and, automatically in response to the detection, causing system control means of the system to allocate bandwidth for transmission of the cells from the ingress to the egress.

19. A method according to claim 18, characterised in that the allocation of bandwidth is preceded by identification of available bandwidth, and in that congestion is signalled to the ingress to cause a reduction in input traffic level.
20. A method according to claim 18 or claim 19, characterised by storing a maximum output cell rate value derived from a control signal received from the system control means, and limiting the transmission of cells to the said switch when the rate at which cells are received at the ingress exceeds the maximum output cell rate value.
21. A method according to claim 20, characterised in that the limiting step includes delaying transmission of cells to the switch in a buffer.
22. A method according to claim 20 or claim 21, characterised in that the maximum output cell rate value is transmitted to a signal source generating the data cells as a feedback signal.
23. A method according to any of claims 20 to 22, characterised by receiving data cells from a plurality of signal sources and characterised in that the limiting step includes transmitting cells received from one of the sources at a rate which is dependent on the number of cells received from one or more of the other sources.
24. A method according to any of claims 18 to 23, characterised by receiving data cells from a plurality of signal sources, and monitoring the incoming cells by reading routing identifiers associated with the cells and updating activity status values associated with the routing

identifiers, the method further including causing the system control means to allocate bandwidth for cells associated with a said identifier upon the transition of an activity status value relating to the said identifier from an active to an inactive status.

25. A method according to any of claims 18 to 24, characterised in that the detecting step includes detecting the rate at which cells are supplied to the ingress from an identified signal source.
26. A method according to claim 25, characterised in that, if the allocated bandwidth is more than that required to transmit the cells at the detected rate, the system control means are caused to allocate less bandwidth to the cells.
27. A method according to claim 18 or claim 24, including reading a bandwidth-identifying portion of the incoming cells and causing the system control means to allocate bandwidth to the cells according to the bandwidth identified by the said portion.
28. A method according to any of claims 18 to 27, characterised by limiting the rate at which incoming cells are accepted onto the network according to a default bandwidth until allocation of bandwidth to the said incoming cells by the system control means has occurred.
29. A method according to any of claims 18 to 28, characterised by counting cells received from respective signal sources thereby to generate charging signals for customer billing.
30. A method according to claim 21 and claim 22, characterised by detecting the filling of a buffer to

a predetermined threshold level and, in response to reaching the said predetermined threshold level, causing retransmission of the maximum output cell rate value for the associated signal source when the
5 buffer has been filled to the threshold level.

31. A dynamic bandwidth controller for a broadband switching system, which system is for transmitting information-carrying asynchronously transferred data
10 cells from an ingress of the system to an egress of the system via at least one switch, wherein the controller comprises means for detecting the said cells received at the ingress, and means for issuing a bandwidth request signal automatically in response
15 to such cell detection to cause bandwidth to be allocated in the system for the transmission of the cells to the egress.

32. A controller according to claim 31, arranged to allow
20 cells to pass through the system while the bandwidth request is being processed in the system and before the said allocation of bandwidth.

33. A controller according to claim 31 or claim 32,
25 characterised by transmission limiting means arranged to store a cell rate indicator value in response to a system control signal, and to limit the onward transmission of cells when the rate at which cells are received at the ingress exceeds a cell rate
30 indicated by the cell rate indicator signal.

1/13




Fig.1.




Fig.2.





Fig. 4.




Fig.5

4/13

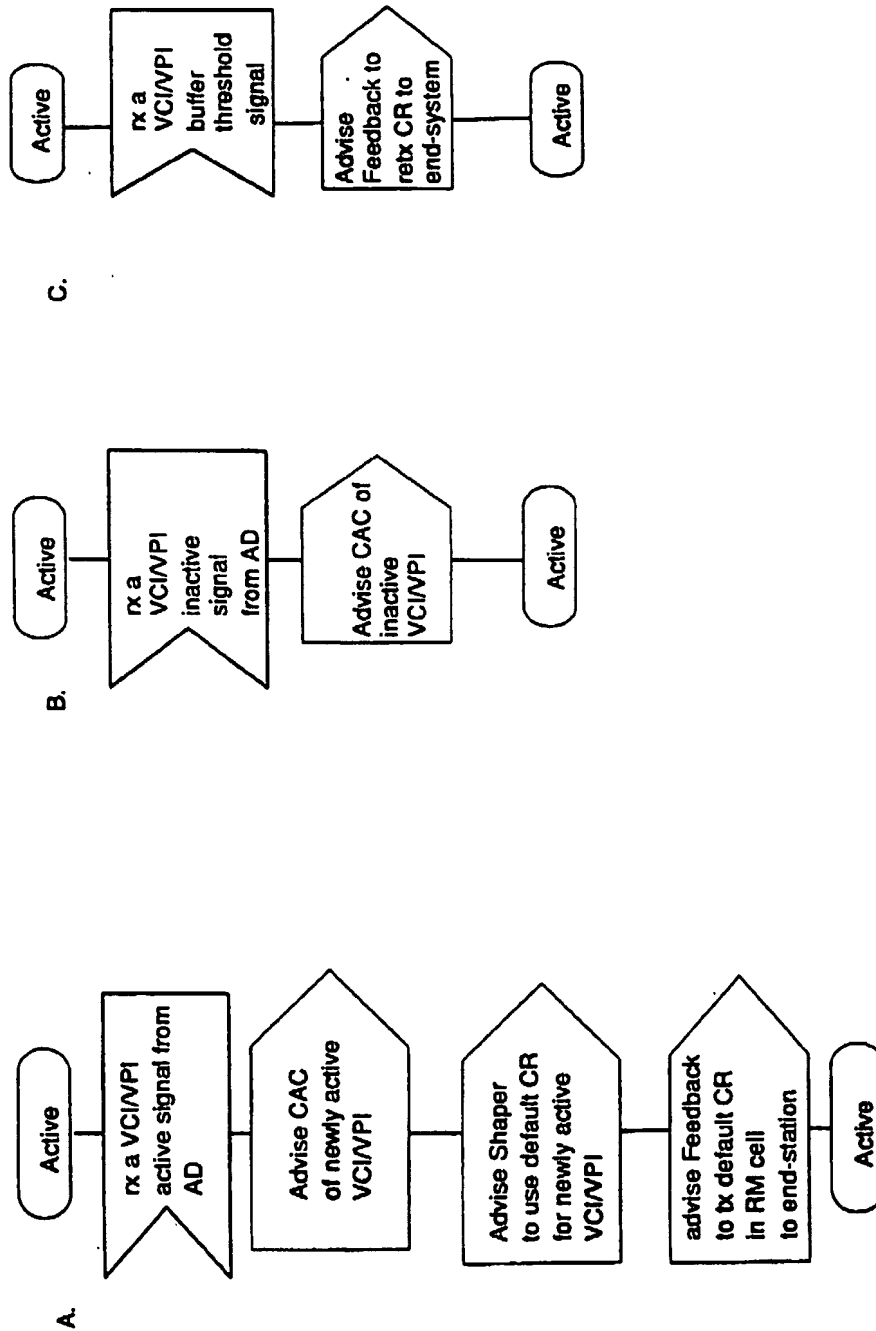


Fig.6-1

5/13

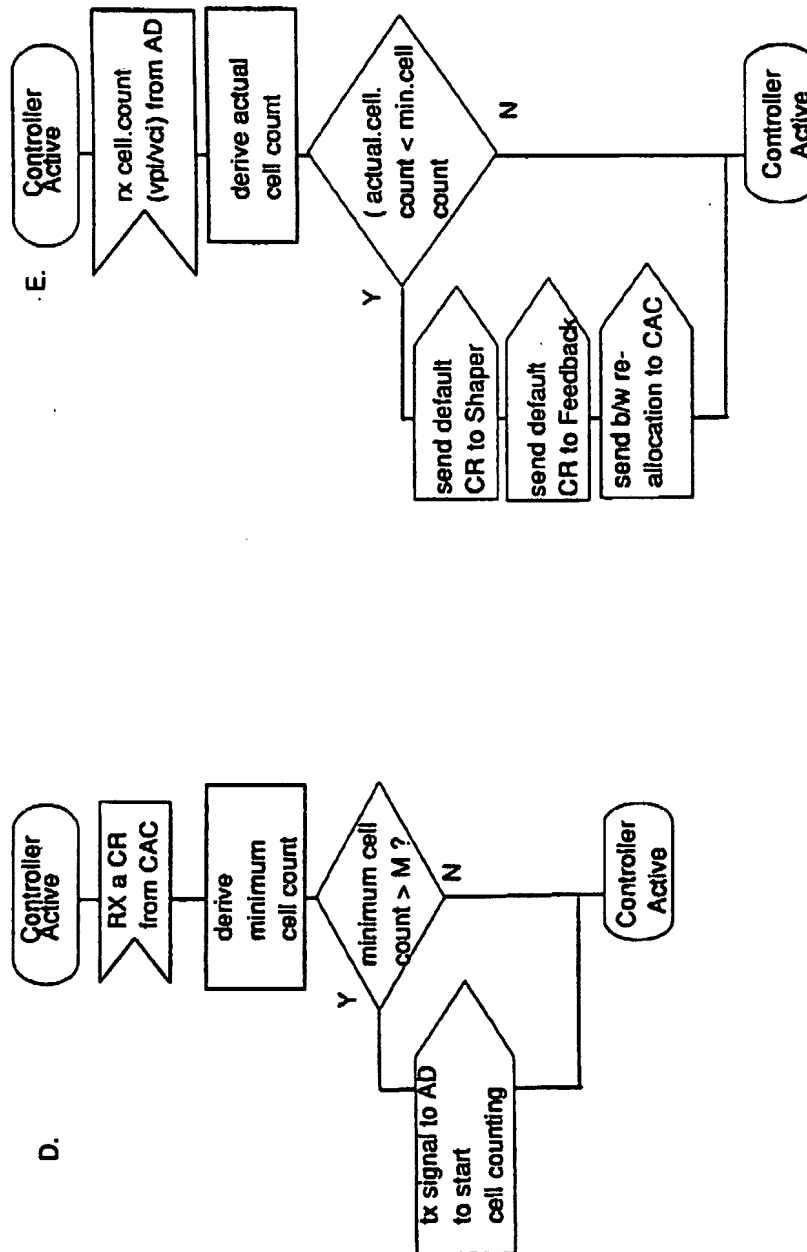





Fig. 6-2



Shaper/multiplexer

VPI/VCI Buffer Fig. 11.

7/13

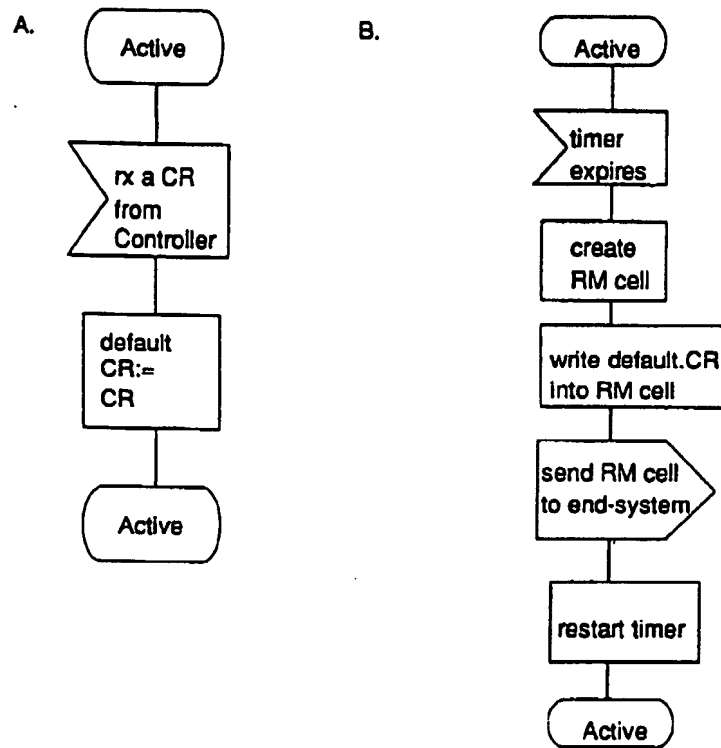


Fig.8

8/13

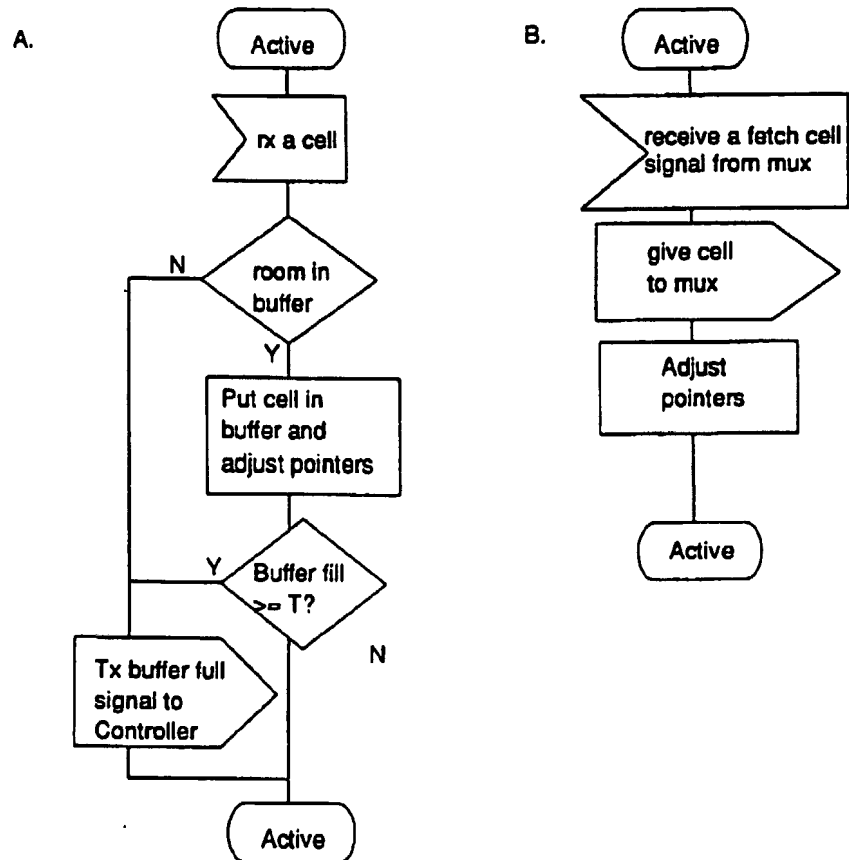


Fig.10

9/13

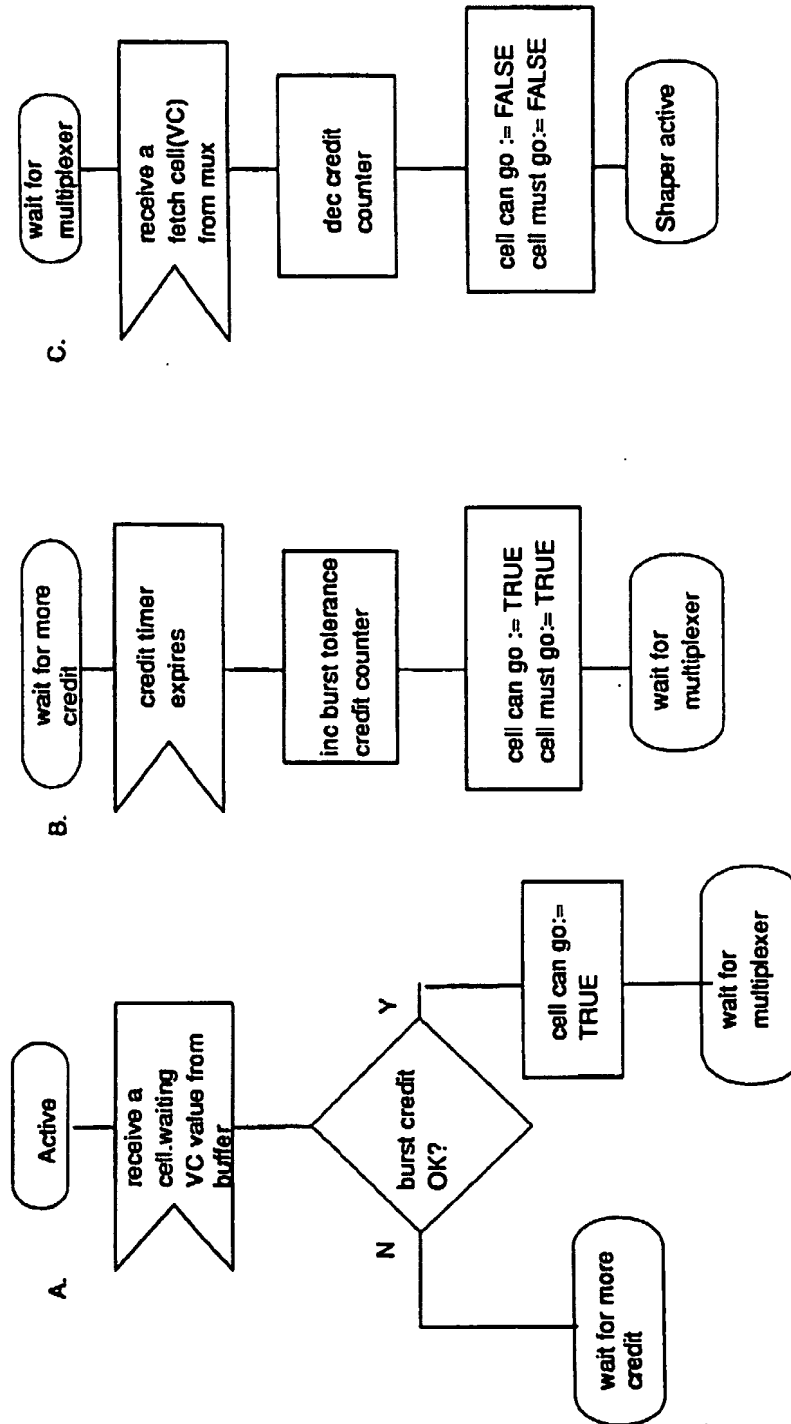


Fig.12-1




Fig.12-2

11/13

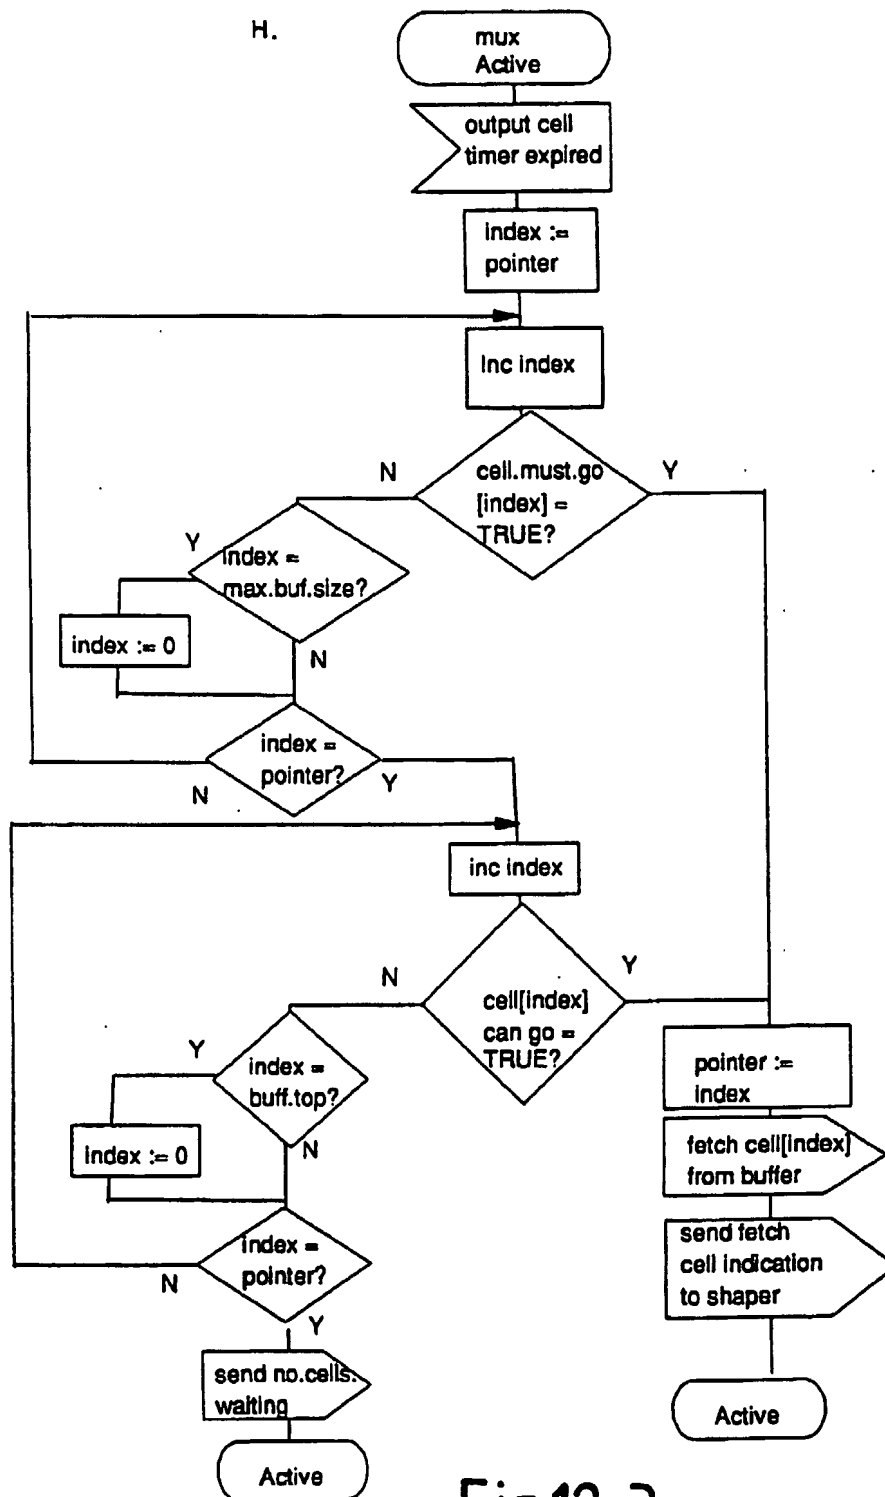


Fig.12-3




Fig.13.

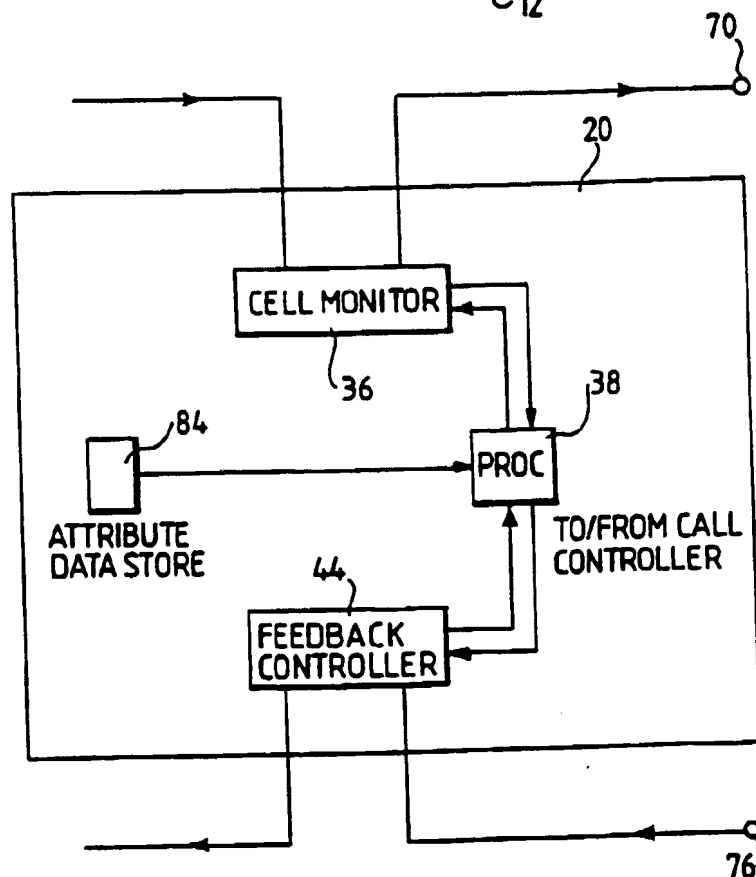


Fig.14.

13/13




Fig.15.

INTERNATIONAL SEARCH REPORT

Intern. Application No
PCT/GB 95/00502

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 H04Q11/04 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 H04L H04Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	IEEE INFOCOM '91, vol. 3, BAL HARBOUR, page 1145 XP 000223441 P. CROCKETT, G. GALLASSI, M. GERLA 'Bandwidth advertising for MAN/ATM connectionless internetting.' see page 1145, right column, line 31 - page 1146, right column, line 13 ---	1, 5, 10-12
A	TRANSACTIONS OF THE INSTITUTE OF ELECTRONICS, INFORMATION AND COMMUNICATIONS ENGINEERS., vol. E74, TOKYO, page 761 XP 000241297 H. SAITO, K. KAWASHIMA, K. SATO. 'Traffic Control Technologies in ATM Networks.' see page 763, left column, line 33 - page 764, right column, line 16 --- -/-	1, 17, 18, 31

☒ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *B* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another claim or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *Z* document member of the same patent family

Date of the actual completion of the international search

6 June 1995

Date of mailing of the international search report

23.06.95

Name and mailing address of the ISA
European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax (+31-70) 340-3016

Authorized officer

Veen, G

INTERNATIONAL SEARCH REPORT

Internat Application No
PCT/GB 95/00502

C(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	INTERNATIONAL SWITCHING SYMPOSIUM, vol. 2, YOKOHAMA, page 7 XP 000337692 D. P. TRANCHIER, P. E. BOYER, Y. M. ROUAUD, J.-Y. MAZEAS 'Fast bandwidth allocation in ATM networks.' see page 7, left column, line 35 - right column, line 45 ---	1,17,18, 31
A	INTERNATIONAL SWITCHING SYMPOSIUM, vol. 1 of 2, 25 October 1992 YOKOHAMA, pages 229-233, XP 000337649 H.J.CHAO 'A General Architecture for Link-Layer Congestion Control in ATM Networks' see page 231, right column, line 14 - line 62 -----	4,7,22